arXiv:1408.2071v1 [cs.DC] 9 Aug 2014

Near-Constant-Time Distributed Algorithms on a Congested Clique*

James W. Hegeman, Sriram V. Pemmaraju, and Vivek B. Sardeshmukh

Department of Computer Science, The University of Iowa, Iowa City, IA 52242 {james-hegeman, sriram-pemmaraju, vivek-sardeshmukh}@uiowa.edu

Abstract. This paper presents constant-time and near-constant-time distributed algorithms for a variety of problems in the congested clique model. We show how to compute a 3-ruling set in expected $O(\log\log\log n)$ rounds and using this, we obtain a constant-approximation to metric facility location, also in expected $O(\log\log\log n)$ rounds. In addition, assuming an input metric space of constant doubling dimension, we obtain constant-round algorithms to compute constant-factor approximations to the minimum spanning tree and the metric facility location problems. These results significantly improve on the running time of the fastest known algorithms for these problems in the congested clique setting.

1 Introduction

The $\mathcal{CONGEST}$ model is a synchronous, message-passing model of distributed computation in which the amount of information that a node can transmit along an incident communication link in one round is restricted to $O(\log n)$ bits, where n is the size of the network [23]. As the name suggests, the $\mathcal{CONGEST}$ model focuses on *congestion* as an obstacle to distributed computation. In this paper, we focus on the design of distributed algorithms in the $\mathcal{CONGEST}$ model on a *clique* communication network; we call this the *congested clique* model. In the congested clique model, all information is nearby, i.e., at most one hop away, and so any difficulty in solving a problem is due to congestion alone.

Let $H = (V, E_H)$ denote the underlying clique communication network. In general, the input to the problems we consider consists of a $|V| \times |V|$ matrix M of edge-attributes and a length-$|V|$ vector of node attributes. M represents edge weights (or distances, or costs) and it is initially distributed among the nodes in V in such a way that each node $v \in V$ knows the corresponding row and column of M. In one typical example, M could simply be the adjacency matrix of a spanning subgraph $G = (V, E)$ of H; in this setting, each node $v \in V$ initially knows all the edges of G incident on it. A number of classical problems in distributed computing, e.g., maximal independent set (MIS), vertex coloring, edge coloring, maximal matching, shortest paths, etc., are well-defined in this setting. However, the difficulty of proving lower bounds in the congested clique model [7] means that it is not clear how quickly one should be able to solve any of these problems in this model. Note that the input G can be quite dense (e.g., have $\Theta(n^2)$ edges) and therefore any reasonably fast algorithm for the problem will have to be "truly" distributed in the sense that it cannot simply rely on shipping off the problem description to a single node for local computation. In this setting, the algorithm of Berns et al. [3,2] that computes a *2-ruling set* of G in expected-$O(\log\log n)$ rounds is worth mentioning. (A *t-ruling set* is defined to be an independent set $I \subseteq V$ such that every node in V is at most t hops in G from some node in I.) In another important class of problems that we study, the input matrix M represents a metric space (V, d); thus each node $v \in V$ initially has knowledge of distances $d(v, w)$ for all $w \in V$. Nodes then need to collaborate to solve a problem such as *minimum spanning tree* (MST) or *metric facility location* (MFL) that are defined on the input metric space. In this setting, the deterministic MST algorithm of Lotker et al. [19] running in $O(\log\log n)$ rounds is worth mentioning.

Thus far the congested clique model has mainly served the theoretical purpose of helping us understand the role of congestion as an obstacle to distributed computation. However, recent papers [15,13] have made connections between congested clique algorithms and algorithms in popular systems of parallel computing such as MapReduce [6] and graph processing systems such as Pregel [20], thus providing a practical motivation for the development of fast algorithms on the congested clique. Specifically, in [13], it is shown that congested

*This work is supported in part by National Science Foundation grant CCF-1318166.

clique algorithms with fairly liberal resource constraints can be efficiently simulated in a MapReduce model of computation [14].

1.1 Main Results

In this paper we present several constant-time or near-constant-time algorithms for fundamental problems in the congested clique setting.

- First, we present an algorithm that computes a 3-ruling set of G in expected $O(\log\log\log n)$ rounds, significantly improving the running time of the 2-ruling set algorithm of Berns et al. [3,2].
- Via a reduction presented in Berns et al. [3,2], this implies an expected $O(\log\log\log n)$-round algorithm for computing an $O(1)$-approximation for MFL. Again, this significantly improves on the running time of the fastest known algorithm for this problem.

Distributed algorithms that run in $O(\log\log n)$ rounds are typically analyzed by showing a doubly-exponential rate of progress; such progress, for example, is achieved if the number of nodes that have "successfully finished" grows by squaring after each iteration. The congested clique algorithms for MST due to Lotker et al. [19] and the above-mentioned MFL algorithm due to Berns et al. [3,2] are both examples of such phenomena. Our algorithm with triply-logarithmic running time, involves new techniques that seem applicable to congested clique algorithms in general. Our result raises the distinct possibility that other problems, e.g., MST, can also be solved in $O(\log\log\log n)$ rounds on a congested clique. In fact, our next set of results represents progress in this direction.

- We show how to solve the MIS problem on a congested clique in *constant* rounds on an input graph G_r induced by the metric space (V, d) in which every pair of nodes at distance at most r (for any $r \geq 0$) are connected by an edge. This result has two implications.
- First, given a metric space (V, d) of constant doubling dimension, we show that a constant-approximation to the MST problem on this metric space can be obtained in *constant* rounds on a congested clique setting.
- An additional implication of the aforementioned MIS result is that it leads to a *constant*-round constant-approximation to MFL in metric spaces of constant doubling dimension on a congested clique.

In order to achieve our results, we use a variety of techniques that balance bandwidth constraints with the need to make rapid progress. We believe that our techniques will have independent utility in any distributed setting in which congestion is a bottleneck.

1.2 Technical Preliminaries

Congested Clique Model. The underlying communication network is a clique $H = (V, E_H)$ of size $n = |V|$. Computation proceeds in synchronous rounds and in each round a node (i) receives all messages sent to it in the previous round, (ii) performs unlimited local computation, and then (iii) sends a, possibly different, message of size $O(\log n)$ to each of the other nodes in the network. We assume that nodes have distinct IDs that can each be represented in $O(\log n)$ bits.

MST and MFL problems. We assume that the input to the MST problem is a metric space (V, d). Initially, each node $v \in V$ knows distances $d(v, w)$ to all nodes $w \in V$. When the algorithm ends, all nodes in V are required to know a spanning tree T of V of minimum weight. (Note that here we take $d(u, v)$ to be the "weight" of edge $\{u, v\}$.) The input to MFL consists of a metric space (V, d) along with *facility opening costs* f_v associated with each node $v \in V$. The goal is to find a subset $F \subseteq V$ of nodes to *open* as facilities so as to minimize the facility opening costs plus connection costs, i.e., $\sum_{v \in F} f_v + \sum_{u \in V} D(u, F)$, where $D(u, F) := \min_{v \in F} d(u, v)$ is the *connection cost* of node u. Initially, each node $v \in V$ knows facility opening cost f_v and distances $d(v, w)$ for all $w \in V$. Facility location is a well-studied problem in operations research [1,4,9] that arises in contexts such as locating hospitals in a city or locating distribution centers in a region. More recently, the facility location problem has been used as an abstraction for the problem of locating resources in a wireless network [10,22] and motivated by this application several distributed approximation algorithms for this problem have been designed [21,11,12].

t-ruling set problem. A *t-ruling set* of a graph $G = (V, E)$ is an independent set $I \subseteq V$ such that every vertex in G is at most t hops from some vertex in I. A t-ruling set, for constant t, is a natural generalization of an MIS and can stand as a proxy for an MIS in many instances. The input to the t-ruling set problem on a congested clique $H = (V, E_H)$ is a spanning subgraph $G = (V, E)$ of the underlying communication network H. Each node $v \in V$ is initially aware of all its neighbors in G. At the end of the t-ruling set algorithm, every node is required to know the identities of all nodes in the computed t-ruling set.

Metric spaces, doubling dimension, and growth-bounded graphs. If $M = (V, d)$ is a metric space then we use $B_M(v, r)$ to denote the set of points $w \in V$ such that $d(v, w) \leq r$. We call $B_M(v, r)$ the *ball of radius r centered at v*. A metric space $M = (V, d)$ has *doubling dimension* ρ if for any $v \in V$ and $r \geq 0$, $B_M(v, r)$ is contained in the union of at most 2^ρ balls $B_M(u, r/2)$, $u \in V$. In this paper, we work with metric spaces with constant doubling dimension, i.e., $\rho = O(1)$. Note that constant-dimensional Euclidean metric spaces are natural examples of metric spaces with constant doubling dimension. In distributed computing literature, metric spaces of constant doubling dimension have been investigated in the context of wireless networks [5,16]. For a graph $G = (V, E)$ and a node $v \in V$, let $B_G(v, r)$ denote the set of all vertices $u \in V$ that are at most r hops from v. A graph $G = (V, E)$ is said to have *bounded growth* (or said to be *growth-bounded*) if the size of any independent set in any ball $B_G(v, r)$, $v \in V$, $r \geq 0$, is bounded by $O(r^c)$ for some constant c. For any metric space (V, d) and $r \geq 0$, the graph $G_r = (V, E_r)$, where $E_r = \{\{u, v\} \in d(u, v) \leq r\}$ is called a *distance-threshold graph*. It is easy to see that if (V, d) has constant doubling dimension then a distance-threshold graph G_r, for any $r \geq 0$, is growth-bounded; this fact will play an important role in our algorithms. For a given metric space (V, d) the *aspect ratio* $\lambda(Y)$ of a subset of points $Y \subseteq V$ is the ratio of maximum of pair-wise distance between points in Y to the minimum of pair-wise distance between points in Y, i.e. $\lambda(Y) = \max\{d(u, v) \mid u, v \in Y\}/\min\{d(u, v) \mid u, v \in Y\}$. The following fact is easy to prove by applying the definition of doubling dimension: if (V, d) is a metric with doubling dimension ρ and $Y \subseteq V$ is a subset of points, then $|Y| \leq 2^{\rho \cdot \lceil \log_2 \lambda(Y) \rceil}$ where $\lambda(Y)$ is the aspect ratio of Y. We refer to this property as the *growth-bounded property* of the metric space (V, d). Distance-threshold graphs and more generally, growth-bounded graphs have attracted attention in the distributed computing community as flexible models of wireless networks [16]. Schneider and Wattenhofer [24] present a deterministic algorithm, running in $O(\log^* n)$ rounds, for computing an MIS on a growth-bounded graph.

Lenzen's routing protocol. A key algorithmic tool that allows us to design constant- and near-constant-time round algorithms is a recent deterministic routing protocol by Lenzen [17] that disseminates a large volume of information on a congested clique in constant rounds. The specific routing problem, called an *Information Distribution Task*, solved by Lenzen's protocol is the following. Each node $i \in V$ is given a set of $n' \leq n$ messages, each of size $O(\log n)$, $\{m_i^1, m_i^2, \ldots, m_i^{n'}\}$, with destinations $d(m_i^j) \in V$, $j \in [n']$. Messages are globally lexicographically ordered by their source i, destination $d(m_i^j)$, and j. Each node is also the destination of at most n messages. Lenzen's routing protocol solves the Information Distribution Task in $O(1)$ rounds.

General Notation. For a subset $S \subseteq V$, $G[S]$ denotes *induced* subgraph of G by set S; thus $G[S] = (S, E')$ where $E' = \{\{u, v\} \mid u, v \in S \text{ and } \{u, v\} \in E\}$. In the context of our MST algorithm we will interpret metric distances $d(u, v)$ as as edge weights; we will use $wt(u, v)$ and $d(u, v)$ interchangeably. Given an edge-weighted graph $G = (V, E)$ and an edge set $E' \subseteq E$, we denote the sum of all edge-weights in E' as $wt(E')$. We use Δ to denote the maximum degree of a graph; sometimes, to avoid ambiguity we use $\Delta(G)$ to denote maximum degree of graph G. All logarithms are assumed to have base 2 unless otherwise specified. We say an event occurs *with high probability* (w.h.p.), if the probability of that event is at least $(1 - 1/n^c)$ for a constant $c \geq 1$.

2 3-Ruling Sets in $O(\log \log \log n)$ Rounds

In this section, we show how nodes in V can use the underlying clique communication network H to compute, in expected-$O(\log \log \log n)$ rounds, a 3-ruling set of an arbitrary spanning subgraph G of H. At a high level, our 3-ruling set algorithm can be viewed as having three steps. In the first step, the graph is decomposed into $O(\log \log n)$ degree-based classes and at the end of this step every node knows the class it belongs to. In the

next subsection, we describe this *degree-decomposition step* and show that it runs in expected $O(\log\log\log n)$ rounds. In the second step, each vertex v of the given graph G joins a set S independently with probability p_v, where p_v depends on v's class as defined in the degree-decomposition step. This *vertex-selection step* yields a set S that will be shown to have two properties: (i) the expected number of edges in the induced subgraph $G[S]$ is $O(n \cdot \text{poly}(\log n))$; and (ii) with high probability, every vertex in G is either in S or has a neighbor in S. Given the degree-decomposition, the vertex-selection step is elementary and requires no communication. In the third step, we use the 2-ruling set algorithm of Berns et al. [3,2]. We show that, on an n-node graph with $O(n \cdot \text{poly}(\log n))$ edges, this algorithm runs in expected-$O(\log\log\log n)$ rounds. We will refer to this algorithm from [3,2] as the *2-ruling set algorithm*. Putting these three steps together yields a 3-ruling set algorithm that runs in $O(\log\log\log n)$ rounds in expectation.

2.1 Degree-Decomposition Step

Let $G = (V, E)$ be an arbitrary graph. Let U_1 be the set of all nodes in G with degrees in the range $[n^{1/2}, n)$. Let V_1 be the remaining nodes, i.e., $V_1 = V \setminus U_1$. Let U_2 be the set of all nodes in V_1 with degrees in $G[V_1]$ belonging to the range $[n^{1/4}, n^{1/2})$. The decomposition continues in this manner until V is partitioned into sets $U_1, U_2, \ldots$. We now provide a more formal description. For $k = 0, 1, 2, \ldots$, let $D_k = n^{1/2^k}$. The D_k's will serve as degree thresholds and will lead to a vertex partition. Let $k^* = \lceil \log\log n \rceil$. Note that $1 < D_{k^*} \leq 2$. Let $V_0 = V$, $G_0 = G$, and $U_1 = \{v \in V_0 \mid \text{degree}_{G_0}(v) \in [D_1, D_0)\}$. For $1 \leq k < k^*$, let

$$V_k = V_{k-1} \setminus U_k, \qquad G_k = G[V_k], \qquad U_{k+1} = \{v \in V_k \mid \text{degree}_{G_k}(v) \in [D_{k+1}, D_k)\}$$

Let $V_{k^*} = V_{k^*-1} \setminus U_{k^*}$, $G_{k^*} = G[V_{k^*}]$, and $U_{k^*+1} = V_{k^*}$. See Figure 1 for an illustration of this decomposition. Let $N_G(v)$ denote the set of neighbors of vertex v in graph G. Here are some easy observations:

(i) For $0 \leq k \leq k^*$, $\Delta(G_k) < D_k$.
(ii) For $1 \leq k \leq k^* + 1$, if $v \in U_k$ then $|N_G(v) \cap V_{k-1}| < D_{k-1}$.
(iii) For $1 \leq k \leq k^* + 1$, if $v \in U_k$ then $|N_G(v) \cap U_j| < D_j$ for $j = 1, 2, \ldots k-1$.

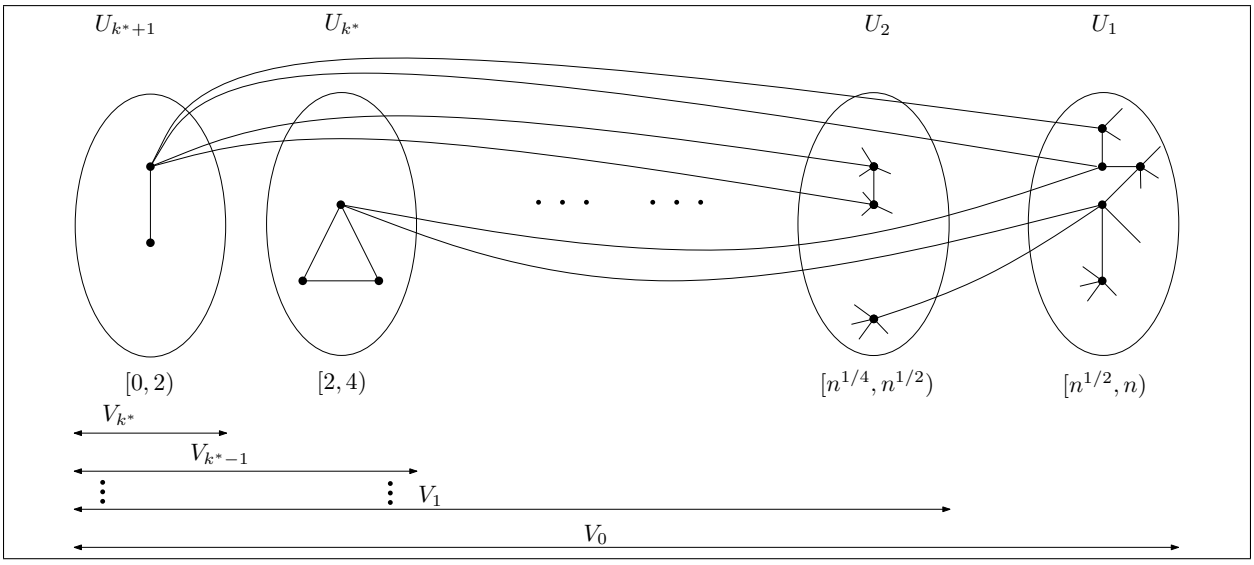

Fig. 1. Degree-Decomposition Step. U_1 is the set of all nodes in G with degrees in the range $[n^{1/2}, n)$ and V_1 is the remaining nodes. U_2 is the set of all nodes in V_1 with degrees in $G[V_1]$ belonging to the range $[n^{1/4}, n^{1/2})$. The decomposition continues in this manner until all nodes belong to some U_k. We use k^* to denote $\lceil \log\log n \rceil$. Assuming that $\log\log n = k^*$, we see that U_k^* is the set of nodes that have degree in $G[V_{k^*-1}]$ in the range $[2, 4)$. Note that a node v that belongs to U_{k+1} could have degree in G that is much larger than $D_k = n^{1/2^k}$.

Now we describe algorithm to compute this degree-decomposition; in particular, we precisely describe how each node v computes an index $k(v) \in [k^*+1]$ such that $v \in U_{k(v)}$. Below, we first describe at a high level a 2-phase approach that we use to compute the index $k(v)$ for each vertex v. Subsequently we will flesh out our approach with necessary details and show that it is correct and can be implemented in $O(\log\log\log n)$ rounds on a congested clique.

Lazy phase: Let $t = \lceil 1 + \log\log\log n \rceil$. The sets $U_1, U_2, \ldots, U_t$ are identified in a leisurely manner, one-by-one, in $O(\log\log\log n)$ rounds. At the end of this phase each vertex $v \in \cup_{i=1}^{t} U_i$ knows the index $k(v) \in [t]$ such that $v \in U_{k(v)}$.

Speedy phase: The set of remaining vertices, namely V_t, induces a graph G_t whose maximum degree is less than

$$D_t \leq n^{1/2^{1+\log\log\log n}} = n^{1/(2\log\log n)}.$$

This upper bound on the maximum degree helps us compute the index values $k(v)$ for the remaining vertices at a faster rate. We first show that each vertex v in G_t can acquire knowledge of the graph induced by the ball $B_{G_t}(v, k^*)$ in $O(\log\log\log n)$ rounds via a fast *ball-growing algorithm*. (Recall that $k^* = \lceil \log\log n \rceil$.) We then show that $G[B_{G_t}(v, k^*)]$ contains enough information for v to determine $k(v) \in [k^* + 1]$ via local computation. Therefore, after each vertex $v \in V_t$ acquires complete knowledge of the radius-k^* ball centered at it, it can locally compute index $k(v)$ and proceed to the vertex-selection step.

We now present the *Lazy-phase algorithm* executed by all vertices $v \in G$.

Algorithm 1 Lazy-phase algorithm at vertex v

```
1. k(v) ← 0
2. for i ← 1 to t do
3.       s(v) ← |{u ∈ N_G(v) | 1 ≤ k(u) < i}|
4.       if degree_G(v) − s(v) ∈ [D_i, D_{i−1}) then
5.             k(v) ← i
6.             Send k(v) to all neighbors
7.             break
```

Lemma 1. *The Lazy-phase algorithm runs in $O(\log\log\log n)$ rounds and at the end of the algorithm, for each vertex $v \in \cup_{j=1}^{t} U_j$, $k(v)$ has a value in $[t]$ such that $v \in U_{k(v)}$. For any vertex $v \notin \cup_{j=1}^{t} U_j$, $k(v)$ is set to 0.*

Proof. Given that the sets $U_1, U_2, \ldots, U_i$ have been determined, and that the members of each are known to every node in the network, each node can locally determine its degree in $G_i = G[V_i]$ and thus determine its membership in U_{i+1}. Each node can then broadcast whether or not it has joined U_{i+1}, thus providing knowledge of U_{i+1} to every node in the network. It follows that the implementation of the Lazy-phase algorithm requires exactly $t = \lceil 1 + \log\log\log n \rceil$ rounds of communication to complete.

We now present the *Speedy-phase algorithm* executed by vertex v. Note that the Speedy-phase algorithm is only executed at vertices v for which $k(v)$ is 0 after the Lazy-phase algorithm. In other words, the Speedy-phase algorithm is only executed at vertices v in G_t, the graph induced by vertices not in $\cup_{j=1}^{t} U_j$. The key idea of the Speedy-phase algorithm is that once each node v in G_t has acquired knowledge of $G_t[B_{G_t}(v, r)]$, then in constant rounds of communication, each node v can "double" its knowledge, i.e., acquire knowledge of $G_t[B_{G_t}(v, 2r)]$. This is done by each node v sending knowledge of $G_t[B_{G_t}(v, r)]$ to all nodes in $B_{G_t}(v, r)$; the key is to establish that this volume of communication can be achieved on a congested clique in constant rounds. This idea has appeared in a slightly different context in [18].

Algorithm 2 Speedy-phase algorithm at vertex v

1. // *Growing the ball* $B_{G_t}(v, k^*)$
2. Each node sends a list of all of its neighbors in G_t to each of its neighbors (in G_t) // *After which each* $v \in V_t$ *knows* $G[B_{G_t}(v, 1)]$
3. **for** $i \leftarrow 0$ **to** $\lceil \log\log\log n \rceil - 1$ **do**
4. Send a description of $G[B_{G_t}(v, 2^i)]$ to all nodes in $B_{G_t}(v, 2^i)$
5. Construct $G[B_{G_t}(v, 2^{i+1})]$ from $G[B_{G_t}(u, 2^i)]$ received from all $u \in B_{G_t}(v, 2^i)$
6. Locally compute $k(v) \in [k^* + 1]$ such that $v \in U_{k(v)}$

Lemma 2. *The Speedy-phase algorithm above runs in* $O(\log\log\log n)$ *rounds in the congested-clique model and when this algorithm completes execution, each vertex* v *in* G_t *knows* $G[B_{G_t}(v, k^*)]$.

Proof. Line 2 of the Speedy-phase algorithm can be completed in a constant number of rounds using Lenzen's routing protocol because each node needs only to send and receive $O(D_t)$ messages to/from $O(D_t)$ neighbors (each message listing a neighbor and destined for a neighbor), as the maximum degree of G_t is less than D_t.

In implementing the Speedy-phase algorithm, the key step is to perform Line 4 in $O(1)$ rounds of communication. If this can be done, then after $O(\log\log\log n)$ rounds, each node v remaining in V_t will have knowledge of its entire neighborhood graph out to a distance of $2^{\lceil \log\log\log n \rceil} \geq \lceil \log\log n \rceil = k^*$ hops away from v.

Since G_t has maximum degree less than D_t, the neighborhood graph $G[B_{G_t}(v, 2^i)]$ can be completely described by listing all $O(D_t^{2^i+1})$ edges. Thus, such a neighborhood can be communicated from v to another node (in particular, to any other node in $B_{G_t}(v, 2^i)$) via $O(D_t^{2^i+1}) = O(n^{(2^i+1)/2^t})$ messages of size $O(\log n)$. Therefore, to perform a given iteration of Line 4 within the Speedy-phase algorithm, each node will need to send (and receive) $O(n^{(2^i+1)/2^t})$ messages (of size $O(\log n)$) to $O(D_t^{2^i}) = O(n^{2^{i-t}})$ other nodes in the network. As above, we can use Lenzen's routing protocol to perform this task in $O(1)$ rounds as long as the total number of messages to be sent (and received) by each node is $O(n)$.

Thus, Line 4 of the Speedy-phase algorithm can be executed in a constant number of rounds if $n^{(2^{i+1}+1)/2^t} = O(n)$; in other words, if $2^{i+1} + 1 \leq 2^t$, or $i \leq t - 2 = \lceil \log\log\log n \rceil - 1$. This lower bound on the maximum value of i that still allows Line 4 to be completed in $O(1)$ rounds is precisely the final index in the for-loop (Line 3). This completes the proof.

Lemma 3. *For any graph* H *and a vertex* v *in* H*, suppose that* v *knows the graph induced by* $B_H(v, k^*)$*. Then* v *can locally compute the index* $k(v) \in [k^* + 1]$ *such that* $v \in U_{k(v)}$.

Proof. The proof is by induction. Whether a vertex u is in U_1 is determined by its degree in H. Since v knows $H[B_H(v, k^*)]$ it can determine via local computation which $u \in B_H(v, k^* - 1)$ belong to U_1 and which don't. As the inductive hypothesis, suppose that for some $i \geq 1$, v has determined for all $u \in B_H(v, k^* - i)$ the following information:

(i) if $u \in \cup_{j=1}^{i} U_j$, then v knows $k(u) \in [i]$ such that $u \in U_{k(u)}$.
(ii) if $u \notin \cup_{j=1}^{i} U_j$, then v knows that $u \notin \cup_{j=1}^{i} U_j$.

Now consider a vertex $u \in B_H(v, k^* - i - 1)$ such that $u \notin \cup_{j=1}^{i} U_j$. In order to determine if $u \in U_{i+1}$, vertex v needs to check if the *residual degree* of u, defined as

$$r(u) := degree_H(u) - |N_H(u) \cap (\cup_{j=1}^{i} U_j)| \tag{1}$$

belongs to the interval $[D_{i+1}, D_i)$. In other words, we need to check that the degree of u after we have deleted all neighbors in $\cup_{j=1}^{i} U_j$ is in the range $[D_{i+1}, D_i)$. Given the information that v knows about all $u \in B(v, k^* - i)$ (by the inductive hypothesis), vertex v can compute the residual degree $r(u)$ for each $u \in B_H(v, k^* - i - 1)$. Therefore for all such u, vertex v can determine if $u \in U_{i+1}$ or not. This completes the inductive step of the proof.

Now since $B_H(v, 0) = \{v\}$, it follows from the above inductive argument that v can determine the index $k(v) \in [k^* + 1]$ such that $v \in U_{k(v)}$.

2.2 Vertex-Selection Step

Algorithm 3 Vertex-Selection Step

if $v \in U_k$ for $k = 1, 2, \ldots, k^*$ **then**
    v is selected with probability $\min\left(\frac{2\log n}{D_k}, 1\right)$
if $v \in U_{k^*+1}$ **then**
    v is selected with probability 1

As mentioned earlier, the vertex-selection step randomly and independently samples nodes in G, with each node v sampled with a probability p_v that depends on the class $U_{k(v)}$ it belongs to. Specifically, if v belongs to U_k then v is independently selected with probability $\min(2\log n/D_k, 1)$. Algorithm 3 shows pseudocode for the vertex-selection step. Let S be the set of vertices that are selected. Let $e(S)$ denote the set of edges in the induced graph $G[S]$.

Lemma 4. $E[|e(S)|] = O(n \cdot \log^2 n \cdot \log\log n)$.

Proof. Consider an arbitrary vertex $v \in V$ and let k, $1 \le k \le k^* + 1$ be such that $v \in U_k$. We will show that the expected number of edges between v and nodes in $\cup_{j \le k} U_j$ is less than $4k \cdot \log^2 n$.

In the graph G, node v has fewer than D_{k-1} neighbors in U_k. Thus, if $1 \le k \le k^*$, the expected number of edges in $e(S)$ between v and nodes in U_k is at most

$$\frac{2\log n}{D_k} \cdot \sum_{u \in N_G(v) \cap U_k} \frac{2\log n}{D_k} < \frac{4\log^2 n \cdot D_{k-1}}{D_k^2} = 4\log^2 n.$$

If $k = k^* + 1$, the number of edges between v and other nodes in U_{k^*+1} is at most 1.

In the graph G, node v has fewer than D_j neighbors in U_j, for $j < k$. Thus, if $1 \le k \le k^*$, the expected number of edges in $e(S)$ between v and nodes in U_j, $j < k$, is at most

$$\frac{2\log n}{D_k} \cdot \sum_{u \in N_G(v) \cap U_j} \frac{2\log n}{D_j} < \frac{4\log^2 n}{D_k} \le 4\log^2 n.$$

If $k = k^* + 1$, the expected number of edges in $e(S)$ between v and nodes in U_j, $j < k$, is

$$1 \cdot \sum_{u \in N_G(v) \cap U_j} \frac{2\log n}{D_j} < 2\log n.$$

Hence, summing over j, the expected total number of edges in $e(S)$ between v and $\cup_{j \le k} U_j$ is less than $4k \cdot \log^2 n$. Using the fact that $k \le 1 + \log\log n$, we see that the expected total number of edges in $e(S)$ between v and $\cup_{j \le k} U_j$ is $O(\log^2 n \cdot \log\log n)$. The result follows.

Lemma 5. *For any $v \in V$,* $\Pr(v$ *is in* S *or* v *has a neighbor in* $S) \ge 1 - 1/n^2$.

Proof. Suppose that $v \in U_k$, for some $1 \le k \le k^*$. Vertex v has at least D_k neighbors in V_{k-1}. Each such neighbor is selected for S with probability at least $\min\{(2\log n)/D_k, 1\}$. If $2\log n \ge D_k$, than any of these neighbors is selected for S with probability 1, so v has a neighbor in S with probability 1. Otherwise, we have

$$\Pr(v \text{ has no neighbor in } S) \le \left(1 - \frac{2\log n}{D_k}\right)^{D_k} \le e^{-2\log n} \le \frac{1}{n^{2.8}}$$

Also, if $v \in U_{k^*+1}$, then v is selected for S with probability 1.

Algorithm 4 2-RulingSet

Input: (G, S)
Output: A 2-ruling set R of $G[S]$

1. $R \leftarrow \emptyset$
2. $m \leftarrow e[G[S]]$ (Each node x broadcasts its degree in $G[S]$ to all others)
3. $q \leftarrow \sqrt{\frac{n}{m}}$
4. **while** $m > 2n$ **do**
5. $\quad T \leftarrow \emptyset$
6. $\quad$ Each $x \in S$ joins T independently with probability q and broadcasts its choice.
7. $\quad$ **if** $e[G[T]] \leq 4n$ **then**
8. $\quad\quad L \leftarrow$LocalMIS$(G[T])$
9. $\quad\quad R \leftarrow R \cup L$
10. $\quad\quad S \leftarrow S \setminus (T \cup N(T))$
11. $\quad m \leftarrow e[G[S]]$
12. $\quad q \leftarrow \sqrt{\frac{n}{m}}$
13. $L \leftarrow$LocalMIS$(G[S])$
14. $R \leftarrow R \cup L$
15. **return** R.

2.3 2-Ruling Set Algorithm

We now briefly describe the 2-ruling set of Berns et al. [3,2] that runs on a congested clique in expected $O(\log\log n)$ rounds. This text is borrowed largely from [3,2] and the reader is advised to consult these papers for a more detailed description. Pseudocode for the algorithm appears in Algorithm 4. The algorithm proceeds in *Iterations* and in each Iteration some number of nodes become inactive and we measure progress by the number of edges remaining in the graph induced by active nodes. In an Iteration i, each active node (S denotes the set of active nodes) joins a "Test" set T independently with probability $q = \sqrt{\frac{n}{m}}$ (Line 6), where m is the number of edges in the graph induced by active nodes. The probability q is set such that the expected number of edges in $G[T]$ is equal to n. If the number of edges in $G[T]$ is no more than $4n$, then we can ship off $G[T]$ to a single node and have that node locally compute an MIS. All of this takes a constant number of rounds and then we delete T and its neighborhood $N(T)$ from the active set S. (Lines 7-10). Because m, the number of edges in $G[S]$, decreases, the probability q rises (Line 12) while still having the expected number of edges in $G[T]$ during the next iteration bounded above by n.

Berns et al. [3,2] have analyzed this algorithm to show that it requires expected $O(\log\log n)$ rounds. We now sketch this analysis to observe that for n-node graphs with $O(n \cdot \text{poly}(\log n))$ edges, this 2-ruling set algorithm requires an expected $O(\log\log\log n)$ rounds.

Lemma 6. *Given an n-vertex graph G with $O(n \cdot \text{poly}(\log n))$ edges, Algorithm 4,* 2-RulingSet *(derived from [3,2]), computes a 2-ruling set of G in expected-$O(\log\log\log n)$ rounds.*

Proof. Algorithm 4, 2-RulingSet computes a 2-ruling set on a subgraph of a congested clique in expected-$O(\log\log n)$ rounds [3,2]. The analysis of this algorithm proceeds by defining $O(\log\log n)$ threshold values $L_k = n^{1+1/2^k}$, for $k = 0, 1, 2, \ldots$, and computing a bound on the expected number of rounds required for the number of edges remaining in the unprocessed portion of the graph to fall from roughly L_i to roughly L_{i+1}. In Lemma 9 of [3], it is proved that this expected number of rounds is uniformly bounded by a constant for every k.

For our use of the 2-ruling set algorithm in the present work, we observe that, if the number of edges in the input subgraph (of the congested clique) is already sufficiently small, then the expected round-complexity of the 2-ruling set algorithm is also much less than would be the case in general. Specifically, we consider a 2-ruling set computation on an input subgraph having $O(n \cdot \text{poly}(\log n))$ edges – in this case, the computation begins having already reached threshold $L_{k'}$, where $n^{1+1/2^{k'}} \approx n \cdot \log^c n$ (for a constant c). More precisely, let $k' = \lfloor \log\log n - \log\log\log n - \log c \rfloor$; then $\frac{1}{2^{k'}} \geq \frac{c \log\log n}{\log n} = \log_n \log^c n$ and $n^{1+1/2^{k'}} \geq n \log^c n$.

Therefore, using the same analysis as occurs in the proof of Theorem 2 in [3] (and in which $\mathcal{T}(k)$ represents the number of iterations necessary to progress from having at most L_{k-1} edges remaining to at most L_k edges remaining), we see that the expected running time (in rounds) of the 2-Ruling Set algorithm applied to an input graph having only $O(n \log^c n)$ edges can be written as

$$\begin{aligned}
\mathbf{E}\left[O(1) + \sum_{k=k'}^{\log\log n} O(\mathcal{T}(k))\right] &= O(1) + \sum_{k=k'}^{\log\log n} O\left(\mathbf{E}[\mathcal{T}(k)]\right) \\
&= O(1) + \sum_{k=\lfloor \log\log n - \log\log\log n - \log c \rfloor}^{\log\log n} O(1) \\
&= O(\log\log\log n)
\end{aligned}$$

which completes the proof.

2.4 Putting it all together

We now combine the algorithm for degree-decomposition step algorithm, the vertex-selection step algorithm, and the 2-ruling set algorithm in order to obtain a 3-ruling set algorithm that runs in $O(\log\log\log n)$ rounds in expectation.

Algorithm 5 3-Ruling Set Algorithm

1. Each node $v \in V$ uses the Lazy-phase and Speedy-phase algorithms to determine the index $k(v) \in [k^* + 1]$ such that $v \in U_{k(v)}$
2. Run the vertex-selection step to compute S
3. $I \leftarrow$ 2-RulingSet($G[S]$)

Lemma 7. *With probability at least $1 - 1/n$, I is a 3-ruling set of G.*

Proof. Consider a vertex $v \in V$. By Lemma 2.2, v has a neighbor in S with probability at least $1 - 1/n^2$. Since I is a 2-ruling set of $G[S]$, there is a node in I at distance at most 3 from v. Thus, with probability at least $1 - 1/n$ we have constructed a 3-ruling set for G.

3 MIS in Growth Bounded Graphs in Constant Rounds

Given a metric space (V, d) with constant doubling dimension, we show in this section how to compute an MIS of a distance-threshold graph $G_r = (V, E_r)$, for any real $r \geq 0$, in a *constant* number of rounds on a congested clique.

3.1 Simulation of the Schneider-Wattenhofer MIS algorithm.

Before we describe our MIS algorithm, we describe an algorithmic tool that will prove quite useful. We know that G_r is growth-bounded and in particular the size of a largest independent set in a ball $B_{G_r}(v, r)$ for any $v \in V$ is $O(r^\rho)$, where ρ is the doubling dimension of (V, d). Schneider and Wattenhofer [24] present a deterministic $O(\log^* n)$-round algorithm to compute an MIS for growth-bounded graphs in the $\mathcal{CONGEST}$ model. Suppose that f is a constant such that the Schneider-Wattenhofer algorithms runs in at most $f \log^* n$ rounds (note that f depends on ρ). We can *simulate* the Schneider-Wattenhofer algorithm in the congested clique model by (i) having each node $v \in V$ grow a ball of radius $f \log^* n$, i.e., gather a description of the induced graph $G[B_{G_r}(v, f \log^* n)]$ and then (ii) having each node v *locally simulate* the Schneider-Wattenhofer algorithm using the description of $G[B_{G_r}(v, f \log^* n)]$. Note that since the Schneider-Wattenhofer algorithm takes at most $f \log^* n$ rounds, it suffices for each node $v \in V$ to know the entire topology of $G[B_{G_r}(v, f \log^* n)]$ to

Algorithm 6 LowDimensionalMIS

Input: $G_r = (V, E_r)$
Output: A maximal independent set $I \subseteq V$ of G_r
1. $P \leftarrow$ ReduceDegree(G_r) // *Phase 1*
2. $(W, Q) \leftarrow$ SampleAndPrune(G_r, P) // *Phase 2*
3. $V' \leftarrow V \setminus (W \cup N(W))$; $R \leftarrow$ SW-MIS(G_r, V') // *Phase 3*
4. $S \leftarrow Q \cup R$; $I \leftarrow$ RulingToMIS(S) // *Phase 4*
5. **return** I

determine if it should join the MIS. The "ball growing" step mentioned above can be implemented by using Lenzen's routing protocol as follows, provided Δ (the maximum degree of G_r) is not too large. Each node v can describe its neighborhood using at most Δ messages of size $O(\log n)$ each. Node v aims to send each of these Δ messages to every node w such that $d(v, w) \leq r \cdot f \log^* n$. In other words, v aims to send messages to all nodes in $B_M(v, r \cdot f \log^* n)$. Since $B_{G_r}(v, f \log^* n) \subseteq B_M(v, r \cdot f \log^* n)$, it follows that the messages sent by v are received by all nodes in $B_{G_r}(v, f \log^* n)$. We now bound the size of $B_M(v, r \cdot f \log^* n)$ as follows. Since M has doubling dimension ρ, the size of any MIS in $B_M(v, r \cdot f \log^* n)$ is $O((\log^* n)^\rho)$ and hence total number of nodes in $B_M(v, r \cdot f \log^* n)$ is $O(\Delta \cdot (\log^* n)^\rho)$. Therefore every node v has $O((\log^* n)^\rho \cdot \Delta^2)$ messages to send, each of size $O(\log n)$. Every node is the receiver of at most $O((\log^* n)^\rho \Delta^2)$ messages by similar arguments. Therefore, if $\Delta = O(\sqrt{n}/(\log^* n)^{\rho/2})$, we can use Lenzen's routing protocol to route these messages in $O(1)$ time. We refer this simulation of the Schneider-Wattenhofer algorithm [24] as Algorithm SW-MIS. The following theorem summarizes this simulation result.

Theorem 1. *If $\Delta(G_r) = O(\sqrt{n}/(\log^* n)^{\rho/2})$ then Algorithm* SW-MIS *computes an MIS of G_r in $O(1)$ rounds on a congested clique.*

3.2 Constant-Round MIS Algorithm

Our MIS algorithm consists of 4 phases. Next we describe, at a high level, what each phase accomplishes.

Phase 1: We compute vertex-subset $P \subseteq V$ such that (i) every vertex in V is at most one hop away from some vertex in P and (ii) $G_r[P]$ has maximum degree bounded above by $c \cdot \sqrt{n}$, for some constant $c > 0$.

Phase 2: We process the graph $G_r[P]$ and compute two subsets W and Q of P such that (i) every vertex in P of degree at least $c \cdot n^{1/4}$ is either in W or has a neighbor in W and (ii) $Q \subseteq W$ is an independent set such that every vertex in W is at most 2 hops from some vertex in Q. Thus, if we delete W and all neighbors of vertices in W what remains is a graph of maximum degree less than $c \cdot n^{1/4}$. Let V' denote the set $P \setminus (W \cup N(W))$. Thus, at the end of Phase 2, Q is a 3-ruling set of $G_r[W \cup N(W)]$ and $\Delta(G_r[V']) < c \cdot n^{1/4}$.

Phase 3: We compute an MIS R of the graph $G_r[V']$ by simply calling SW-MIS.

Phase 4: Since Q is a 3-ruling set of $G_r[W \cup N(W)]$ and R is an MIS of $G_r[V']$, we see that $Q \cup R$ is a 3-ruling set of $G_r[P]$ and thus a 4-ruling set of G_r. In the final phase, we start with the 4-ruling set $Q \cup R$ and expand this into an MIS I of G_r.

Phase 2 is randomized and runs in constant rounds w.h.p. The remaining phases are deterministic and run in constant rounds each. Algorithm LowDimensionalMIS summarizes our algorithm. We now describe each phase in more detail.

3.3 Phase 1: Reduce Degree to $O(\sqrt{n})$

Algorithm ReduceDegree describes Phase 1 of our algorithm. The algorithm consists of arbitrarily partitioning the vertex-set of G_r into $\sqrt{n}$ groups of size (roughly) $\sqrt{n}$ each and then separately and in parallel computing an MIS of each part. Since each part has $\sqrt{n}$ vertices, each part induces a subgraph with at most n edges and therefore each such subgraph can be shipped off to a distinct node and MIS on each subgraph can be computed locally. (The subroutine LocalMIS in Line 4 refers to an unspecified MIS algorithm that is executed locally at a node.) Using the fact that G_r is growth-bounded, we show that the union of all the MIS sets (set P, Line 5) induces a graph with maximum degree bounded by $c \cdot \sqrt{n}$ for some constant c. Also, we show that Phase 1 runs in constant rounds (Lemma 8).

Algorithm 7 ReduceDegree (Phase 1)

Input: $G_r = (V, E_r)$
Output: $P \subseteq V$ such that (i) $V = P \cup N(P)$ and (ii) $\Delta(G_r[P]) < c \cdot \sqrt{n}$ for some constant $c > 0$.
1. Partition V (arbitrarily) into $\lceil \sqrt{n} \rceil$ subsets: $V_1, V_2, \ldots V_{\lceil \sqrt{n} \rceil}$, each of size at most $\sqrt{n}$
2. **for all** $i \leftarrow 1$ **to** $\lceil \sqrt{n} \rceil$ **in parallel do**
3. Send $G_r[V_i]$ to a vertex v_i with lowest ID in V_i
4. Vertex v_i executes $P_i \leftarrow$ LocalMIS($G_r[V_i]$)
5. $P \leftarrow \cup_{i=1}^{\lceil \sqrt{n} \rceil} P_i$
6. **return** P

Algorithm 8 SampleAndPrune (Phase 2)

Input: (G_r, P)
Output: (W, Q), $W \subseteq P$ such that $\{v \in P \mid \text{degree}_{G_r[P]}(v) \geq n^{1/4}\} \subseteq W \cup N(W)$; independent set $Q \subseteq W$ such that Q is a 2-ruling set of $G_r[W]$.
1. **for all** $v \in P$ **in parallel do**
2. Vertex $v \in P$ adds itself to W_i with probability $1/n^{1/4}$ for $i = 1, 2, \ldots, \lceil 2 \cdot \log n \rceil$.
3. $W \leftarrow \cup_{i=1}^{\lceil 2 \log n \rceil} W_i$
4. **for all** $i \leftarrow 1$ **to** $\lceil 2 \log n \rceil$ **in parallel do**
5. Send $G_r[W_i]$ to a vertex w_i, where w_i is the vertex of rank i in the sequence of vertices in V sorted by increasing ID
6. Vertex w_i executes $X_i \leftarrow$ LocalMIS($G_r[W_i]$)
7. $Q \leftarrow$ SW-MIS($G_r[\cup_{i=1}^{\lceil 2 \log n \rceil} X_i]$)
8. **return** (W, Q)

Lemma 8. *Algorithm* ReduceDegree *completes in $O(1)$ rounds and returns a set P such that $\Delta(G_r[P]) < c \cdot n^{1/2}$ for some constant $c > 0$ (that depends on the doubling dimension of the underlying space).*

Proof. Algorithm ReduceDegree starts by arbitrarily partitioning V into $\lceil \sqrt{n} \rceil$ disjoint subsets $V_1, \ldots, V_{\lceil \sqrt{n} \rceil}$ each of size at most $\sqrt{n}$ which can be done in $O(1)$ rounds easily. Since $|V_i| \leq \sqrt{n}$, $G_r[V_i]$ contains at most n edges, for any $i \in \lceil \sqrt{n} \rceil$. Using Lenzen's routing protocol, all knowledge of $G_r[V_i]$ can be shipped off to a designated vertex v_i in V_i (e.g., vertex with smallest ID in V_i) in $O(1)$ rounds. The vertex v_i then computes an MIS P_i of $G_r[V_i]$ locally as shown in Line 4 of Algorithm ReduceDegree. Finally, v_i informs vertices in P_i of their selection into the MIS. The union of the P_i's, denoted P, is returned by the algorithm. This discussion shows that Algorithm ReduceDegree completes in $O(1)$ rounds.

Consider a vertex $u \in P_i$ for some $i \in [\lceil \sqrt{n} \rceil]$. In $G_r[P]$, vertex u cannot have neighbors in P_i since P_i is an independent set in $G_r[P]$. Consider a set P_j, $j \neq i$. The distance between any two vertices in $N(u) \cap P_j$ must be more than r (these nodes are independent) and it must be at most $2r$ (by the triangle inequality). Since the underlying metric space has doubling dimension ρ, it follows that $|N(u) \cap P_j| \leq 2^\rho$. Hence the degree of u in $G_r[P]$ is bounded above by $2^\rho \cdot (\lceil \sqrt{n} \rceil - 1)$. The result follows.

3.4 Phase 2: Sample and Prune

Algorithm SampleAndPrune implements Phase 2 of our MIS algorithm. It takes the induced subgraph $G_r[P]$ as input and starts by computing a set $W \subseteq P$ using a simple random sampling approach. Specifically, for each $i = 1, 2, \ldots, \lceil 2 \cdot \log n \rceil$, each vertex in P simply adds itself to a set W_i independently, with probability $1/n^{1/4}$. We start by proving a useful property of W.

Lemma 9. *Every node u with degree at least $n^{1/4}$ in $G_r[P]$ has a neighbor in W with probability at least $1 - \frac{1}{n^2}$.*

Proof. Let $u \in P$ be a node with degree at least $n^{1/4}$ in $G_r[P]$. For any neighbor v of u, $\Pr(v \notin W) \leq \left(1 - \frac{1}{n^{1/4}}\right)^{\lceil 2 \log n \rceil}$. Therefore the probability that no neighbor of u is in W is at most $\left(1 - \frac{1}{n^{1/4}}\right)^{\lceil 2 \log n \rceil \cdot n^{1/4}}$. This is bounded above $e^{-\lceil 2 \log n \rceil}$, which is bounded above by $1/n^2$.

After using random sampling to compute W, Algorithm SampleAndPrune then "prunes" W in constant rounds to construct a subset $Q \subseteq W$ such that Q is a 2-ruling set of W. In the rest of this subsection we prove that Algorithm SampleAndPrune does behave as claimed here.

Lemma 10. *The number of edges in $G_r[W_i]$ is $O(n)$ w.h.p., for each $i = 1, 2, \ldots, \lceil 2 \log n \rceil$.*

Proof. We first bound the size of the set W_i and the maximum degree of $G_r[W_i]$ for any $i = 1, 2, \ldots, \lceil 2 \log n \rceil$. Observe that $\mathbf{E}[|W_i|] = n^{3/4}$ and since nodes join W_i independently, an application of Chernoff's bound [8] yields $\Pr(|W_i| \leq 6n^{3/4}) \geq 1 - \frac{1}{n^2}$. To bound $\Delta(G_r[W_i])$ we use the fact that degree of any node in $G_r[P]$ is at most $\sqrt{n}$ and therefore the expected degree of any node in $G_r[W_i]$ is at most $n^{1/4}$. Another application of Chernoff's bound yields $\Pr(\text{degree}_{G_r[W_i]}(v) \leq 6n^{1/4}) \geq 1 - \frac{1}{n^2}$ for each node v. Using the union bound over all nodes $v \in W_i$ yields that with probability at least $1 - \frac{1}{n}$ every node in $G_r[W_i]$ has degree at most $6n^{1/4}$. Hence, with high probability, the number of edges in $G[W_i]$ is at most $36n$.

Lemma 11. *The set $X := \cup_{i=1}^{\lceil 2 \log n \rceil} X_i \subseteq P$ is computed in constant rounds w.h.p. in Lines 4-6 of Algorithm* SampleAndPrune*. Furthermore, Every vertex in W is at most one hop away from some vertex in X.*

Proof. We argue that Line 5 can be implemented in $O(1)$ rounds w.h.p. By Lemma 10, each node has to send at most $O(n^{1/4})$ messages to w_i and w.h.p. eachw_i receives at most $O(n)$ messages. Therefore by Lenzen's routing protocol Line 5 takes $O(1)$ rounds. To repeat this for each $i = 1, 2, \ldots, \lceil 2 \log n \rceil$ in parallel, every node has to send at the most $\lceil 2 \log n \rceil \cdot n^{1/4}$ messages. Since w_i's are distinct no w_i needs to receive more than $O(n)$ messages.

Each $v \in W$ belongs to W_i for some i and is therefore at most one hop from some vertex in X_i.

Lemma 12. *W.h.p. it takes constant number of rounds to compute Q. Furthermore, Q is a 2-ruling set of $G_r[W]$.*

Proof. Consider a node $v \in \cup_{i=1}^{\lceil 2 \log n \rceil} X_i$. Since each X_i is an independent set, by using the growth-bounded property of $G_r[X_i]$, we see that the number of neighbors of v in X_i is bounded above by a constant. Hence, the maximum degree in $G_r\left[\cup_{i=1}^{\lceil 2 \log n \rceil} X_i\right]$ is $O(\log n)$. Since the maximum degree of this growth-bounded graph is $O(\log n)$, by Theorem 1 an MIS of this graph can be computed in constant rounds by using SW-MIS.

A node $v \in W$ belongs to some W_i and is therefore at most one hop from some node in X_i. Also, every node in every X_i is at most one hop from some node in Q. Also, Q is independent and therefore Q is a 2-ruling set of $G_r[W]$.

3.5 Phase 4: Ruling Set to MIS

Algorithm RulingToMIS implements Phase 4 of our MIS algorithm. The algorithm takes as input the graph G_r and the vertex subset $S = Q \cup R$ where Q and R are the outputs of Phase 2 and Phase 3, respectively. Note that Lemma 12 implies that S is a 4-ruling set of G_r. This property is used to cover G_r with balls of radius $4r$, centered at members of S.

Consider the graph $G_{9r} = (V, E_{9r})$ where $E_{9r} = \{\{u, v\} \mid u, v \in V \text{ and } d(u, v) \leq 9r\}$. In Lemma 13 we prove a constant upper bound on the maximum degree $\Delta(G_{9r}[S])$. This allows us to compute a proper vertex coloring of $G_{9r}[S]$ using a constant number of colors. This coloring guides the rest of the algorithm, providing a schedule for processing the vertices in the aforementioned balls centered at vertices in S. For each color i, the algorithm processes all vertices in S colored i in parallel. For each vertex $v \in S$ colored i, let B_v denote the subset of $B(v, 4r)$ of vertices still "active". The algorithm computes an MIS of the induced subgraph $G_r[B_v]$; this computation occurs in parallel for each v colored i. Since the vertex coloring is with respect to G_{9r}, two balls B_v and $B_{v'}$ that are processed in parallel do not intersect and in fact are not even connected by an edge. Thus processing in parallel all of the balls B_v for v colored i has no untoward consequences. We note that due to the growth bounded property, every independent set of $G_r[B_v]$ has a constant number of vertices. Hence, we can use a simple sequential algorithm to compute an MIS of $G_r[B_v]$ – repeatedly each vertex with smallest ID in its neighborhood joins the MIS and the graph is updated. We call this MIS algorithm SequentialMIS and use it in Line 9 in Algorithm RulingToMIS. Since every vertex in V is at distance at most $4r$ from some vertex in S, every vertex in V is is some ball B_v and is eventually processed.

Algorithm 9 RulingToMIS (Phase 4)

Input: $(G_r, S = Q \cup R)$
Output: A maximal independent set $I \subseteq V$ of G_r
1. $E_{9r} \leftarrow \{\{u, v\} \mid u, v \in S \text{ and } d(u, v) \leq 9r\}$
2. $G_{9r}[S] \leftarrow (S, E_{9r})$
3. Send $G_{9r}[S]$ to a vertex v^* with lowest ID in S
4. Vertex v^* executes $\Psi \leftarrow \text{LocalColoring}(G_{9r}[S])$ with color pallet $\{1, 2, \ldots, \gamma + 1\}$. Here γ is the constant from Lemma 13.
5. $V' \leftarrow V$
6. **for** $i = 1$ **to** $i = \gamma + 1$ **do**
7. **for all** $v \in S$ such that $\Psi(v) = i$ **in parallel do**
8. $B_v \leftarrow \{u \mid u \in V' \text{ and } d(u, v) \leq 4r\}$
9. $I_v \leftarrow \text{SequentialMIS}(G_{9r}[B_v])$
10. $V' \leftarrow V' \setminus \left(\cup_{v \in S \wedge \Psi(v) = i} (N(I_v))\right)$
11. $I \leftarrow \cup_{v \in S} I_v$
12. **return** I

Lemma 13. *$\Delta(G_{9r}[S]) \leq \gamma$, where γ is a constant.*

Proof. Consider any node $v \in S$ and neighbors $N_{G_{9r}}(v)$ of v in $G_{9r}[S]$. By the triangle inequality, any pair of nodes in $N_{G_{9r}}(v)$ are at most distance $18r$ apart and by Lemma 12, at least distance r apart. Hence $N_{G_{9r}}(v) \cup v$ has a constant aspect ratio and by the growth-bounded property, we have $|N_{G_{9r}}(v) \cup v| \leq 18^\rho = \gamma$.

Lemma 14. *Algorithm* RulingToMIS *executes in a constant number of rounds.*

Proof. Since the maximum degree of $G_{9r}[S]$ is a constant, the entire description of $G_{9r}[S]$ can be shipped to a designated vertex v^* (e.g., a vertex with the smallest ID) using Lenzen's routing protocol in $O(1)$ rounds. Then v^* can compute a coloring of $G_{9r}[S]$ such that no two adjacent vertices have the same color. Notice that the maximum degree of $G_{9r}[S]$ is bounded above by γ, hence $\gamma + 1$ colors are sufficient.

The constant upper bound on the size of the color palette implies that the **for**-loop starting in Line 6 executes a constant number of iterations. In each iteration i, all nodes $v \in S$ colored i are processed. Specifically, an MIS of $G_r[B_v]$ is computed and since the size of every independent set in $G_r[B_v]$ is bounded above by a constant (by appealing to the growth-bounded property), Algorithm SequentialMIS terminates in constant rounds. Hence, each iteration of the outer-**for**-loop takes a constant number of rounds of communication.

Lemma 15. *The set I computed by Algorithm* RulingToMIS *is an MIS of G_r.*

Proof. First we show that I is an independent set by contradiction. Suppose that for some $p, q \in I$, p and q are adjacent in G_r. Then it must be the case that both p and q were selected in the same iteration of the outer-**for**-loop; otherwise, the selection of one of the two nodes would render the other unavailable for selection. If p and q are selected in the same outer-**for**-loop iteration, it must be the case that $p \in B_v$ and $q \in B_{v'}$ where $v \neq v'$, but v and v' have the same color. Since $d(p, v) \leq 4r$, $d(q, v') \leq 4r$, and $d(p, q) \leq r$, using the triangle inequality we see that $d(v, v') \leq 9r$. But, if this is the case then there is an edge between v and v' in $G_{9r}[S]$ and these two vertices would not have the same color, contradicting our earlier conclusion that v and v' have the same color.

We now prove that I is maximal. Since S is a 4-ruling set of G_r, every node $u \in V$ is in $B(v, 4r)$ for some $v \in S$. Suppose that v is colored i and therefore B_v is processed in iteration i of the outer-**for**-loop. If $u \in B_v$ then Algorithm SequentialMIS will either pick u or a neighbor to join the MIS. Otherwise, if $u \notin B_v$ then it must be the case that in an earlier iteration of the outer-**for**-loop, either u or a neighbor were selected to be in the MIS.

4 Constant-Approximation to MST in Constant Rounds

For a metric space (V, d), define a *metric graph* $G = (V, E)$ as the clique on set V with each edge $\{u, v\}$ having weight $d(u, v)$. In this section we present a constant-round algorithm for computing a constant-factor

Algorithm 10 MST-Approximation

Input: A metric graph $G = (V, E)$ on metric space (V, d)
Output: A tree $\hat{\mathcal{T}}$ such that $wt(\hat{\mathcal{T}}) = O\left(wt\left(MST\left(G\right)\right)\right)$

1. $d_m = \max\{d(u,v) \mid \{u,v\} \in E\}$
2. $E_\ell \leftarrow \left\{\{u,v\} \mid d(u,v) \leq \frac{d_m}{n^3}\right\}$ // *Processing light edges*
3. $S \leftarrow$ ComputeMIS$(G[E_0])$ where $E_0 \leftarrow \left\{\{u,v\} \mid d(u,v) \leq \frac{d_m}{n^2}\right\}$
4. $\hat{E}_\ell \leftarrow \left\{\{u,v\} \mid u \in S \text{ and } d(u,v) \leq \frac{2 \cdot d_m}{n^2}\right\}$
5. $E_\hbar \leftarrow \left\{\{u,v\} \mid d(u,v) > \frac{d_m}{n^3}\right\}$ // *Processing heavy edges*
6. $h \leftarrow \left\lceil \frac{3 \log n}{\log c_1} \right\rceil$; $r_0 \leftarrow \frac{d_m}{c_1^h}$
7. **for** $i = 1$ **to** h **in parallel do**
8. $\quad r_i \leftarrow (c_1)^i \cdot r_0$
9. $\quad E_i \leftarrow \{\{u,v\} \mid d(u,v) \leq r_i\}$
10. $\quad V_i \leftarrow$ ComputeMIS$(G[E_i])$
11. $\quad \hat{E}_i \leftarrow \{\{u,v\} \mid u,v \in V_i \text{ and } d(u,v) \leq c_2 \cdot r_i\}$
12. $\hat{E}_\hbar \leftarrow \cup_{i=1}^{h} \hat{E}_i$; $\hat{E} \leftarrow \hat{E}_\ell \cup \hat{E}_\hbar$
13. **return** MST-Sparse$(G[\hat{E}])$

approximation of an MST of given metric graph $G = (V, E)$ with constant doubling dimension. We require that at the end of the MST algorithm, each node in V know the entire spanning tree. Our overall approach is as follows. We start by showing how to "sparsify" G and construct a spanning subgraph $\hat{G} = (V, \hat{E})$, $\hat{E} \subseteq E$, such that $wt(MST(\hat{G})) = O(wt(MST(G)))$. Thus computing an MST on $\hat{G}$ yields an $O(1)$-approximation to an MST on G. The sparsification is achieved via the construction of a collection of maximal independent sets (MIS) *in parallel* on different distance-threshold subgraphs of G. Thus we have reduced the problem of constructing a constant-approximation of an MST on the metric graph G to two problems: (i) the MIS problem on distance-threshold graphs and (ii) the problem of computing an MST of a sparse graph $\hat{G}$. Using the fact that the underlying metric space (V, d) has constant doubling dimension, we show that $\hat{G}$ has linear (in $|V|$) number of edges. As a result, problem (ii) can be easily solved in constant number of rounds by simply shipping $\hat{G}$ to a single node for local MST computation. In Section 3, we have already shown how to compute an MIS of a distance-threshold graph in a constant doubling dimensional space on a congested clique in constant number of rounds. Finally, we show that due to the particular bandwidth usage of our MIS algorithm, we can run all of the requisite MIS computations in parallel in constant rounds.

4.1 MST Algorithm

We now present our algorithm in detail; the reader is encouraged to follow along the pseudocode in Algorithm 10. We partition the edge set E of the metric graph into two subsets E_ℓ (*light* edges) and $E_\hbar$ (*heavy* edges) as follows. Let $d_m = \max\{d(u,v) \mid \{u,v\} \in E\}$ denote the diameter of the metric space [1]. Define $E_\ell = \left\{\{u,v\} \mid d(u,v) \leq d_m/n^3\right\}$ and $E_\hbar = E \setminus E_\ell$. We deal with these two subsets E_ℓ and $E_\hbar$ separately.

First consider the set of light edges E_ℓ and note that $G[E_\ell]$ may have several components. We would like to select an edge set $\hat{E}_\ell$ such that (i) any pair of vertices that are in the same connected component in $G[E_\ell]$ are also in the same connected component in $G[\hat{E}_\ell]$, and (ii) $wt(\hat{E}_\ell) = O(wt(MST(G)))$. (Note that one can define $\hat{E}_\ell = E_\ell$ to have these two properties but we want to "sparsify" E_ℓ, ideally we would like to have $|\hat{E}_\ell| = O(n)$ and we show this for metric with constant doubling dimension.) The algorithm for selecting $\hat{E}_\ell$ is as follows. Let S be an MIS of the distance-threshold graph G_r, where $r = d_m/n^2$. (This MIS computation is not on graph induced by E_ℓ, notice the r. This is done to obtain certain properties of $\hat{E}_\ell$ described above.) Define $\hat{E}_\ell = \left\{\{u,v\} \mid u \in S \text{ and } d(u,v) \leq 2 \cdot d_m/n^2\right\}$. Note that $\hat{E}_\ell$ may not be a subset of E_ℓ.

Now we consider the set $E_\hbar$ of heavy edges. Let $c_1 > 1$ be a constant. Let h be the smallest positive integer such that $c_1^h \geq n^3$. Observe that $h = \left\lceil \frac{3 \log n}{\log c_1} \right\rceil$. Let $r_0 = d_m/c_1^h$ (note that for any heavy edge $\{u,v\}$, $d(u,v) > r_0$) and let $r_i = c_1 \cdot r_{i-1}$, for $i > 0$. We construct $\hat{E}_\hbar$ in *layers* as follows. Let $V_0 = V$ and V_i

[1] If the size of the encoding of distances is more than $O(\log n)$ bits then it is suffices to know only most-significant $\log n$-bits of encoding of d_m to act as "proxy" for d_m which will only increase the approximation factor by a constant.

for $0 < i \leq h$ is an MIS of the subgraph $G[E_i]$ where $E_i = \{\{u,v\} \mid d(u,v) \leq r_i\}$. Let $c_2 > c_1 + 2$ be a constant. Define $\hat{E}_i$, the edge set at the layer i as: $\hat{E}_i = \{\{u,v\} \mid u,v \in V_i \text{ and } d(u,v) \leq c_2 \cdot r_i\}$. We define $\hat{E}_{\hbar} = \cup_{i=1}^{h} \hat{E}_i$ and $\hat{E} = \hat{E}_{\hbar} \cup \hat{E}_{\ell}$. A key feature of our algorithm is that a layer $\hat{E}_i$ does not depend on other layers and therefore these layers can be constructed in parallel. We then call an as-yet-unspecified algorithm called MST-Sparse that quickly computes an exact MST of $\hat{G} = G[\hat{E}]$ in the congested clique model.

In the analysis that follows, we separately analyze the processing of light edges and heavy edges. We first show the *constant-approximation property* of $\hat{G}$ which doesn't require metric to be of constant doubling dimension. Later we show if the underlying metric has constant doubling dimension then Algorithm 10 runs in constant rounds w.h.p..

4.2 Constant-Approximation Property

Let $\mathcal{T}$ be an MST of graph $G = (V, E)$. Let $\hat{\mathcal{T}}$ be a MST of the graph $\hat{G} = (V, \hat{E})$. We now prove that $wt(\hat{\mathcal{T}}) = O(wt(\mathcal{T}))$. First we claim that the connectivity that edges in E_ℓ (i.e., the light edges) provide is preserved by the edges selected into $\hat{E}_\ell$ (Lemma 16) and the total weight of these selected edges is not too high (Lemma 17). Later we prove a similar claim for heavy edges (Lemma 18).

Lemma 16. *For any vertices s and t in V, if there is a s-t path in $G[E_\ell]$ then there exists an s-t path in $G[\hat{E}_\ell]$.*

Proof. Consider an edge $\{u,v\} \in E_\ell$. If $\{u,v\} \in \hat{E}_\ell$ then we are done. If $\{u,v\} \notin \hat{E}_\ell$ then we show that there exists a vertex w such that $\{u,w\}, \{v,w\} \in \hat{E}_\ell$. Since $\{u,v\} \in E_\ell$, $d(u,v) \leq d_m/n^3$. Furthermore, since $\{u,v\} \notin \hat{E}_\ell$ it means both u and v are not in S, an MIS of G_r, $r = d_m/n^2$. Hence there is a vertex $w \in S$ such that $d(u,w) \leq d_m/n^2$. By the definition of $\hat{E}_\ell$, $\{u,w\} \in \hat{E}_\ell$. By the triangle inequality, we have $d(v,w) \leq d_m/n^2 + d_m/n^3$ which implies $\{v,w\} \in \hat{E}_\ell$. The lemma follows by repeatedly applying above result to each edge of the given s-t path.

Lemma 17. $wt(\hat{E}_\ell) = O(wt(\mathcal{T}))$.

Proof. The weight of each edge in $\hat{E}$ is at most $2d_m/n^2$ and since there are at most n^2 edges in $\hat{E}_\ell$ (trivially), we see that $wt(\hat{E}_\ell) = O(d_m)$. We obtain the lemma by using the fact that the total weight of any spanning tree is bounded below by d_m.

Consider an edge $\{u,v\} \in E(\mathcal{T})$. Let $C(u)$ and $C(v)$ be the components containing u and v respectively in the graph $\mathcal{T} \setminus \{u,v\}$.

Lemma 18. *If $\{u,v\} \in E(\mathcal{T}) \cap E_{\hbar}$ then there exists an edge $\{u',v'\} \in \hat{E}$ such that (i) $d(u',v') \leq c_2 \cdot d(u,v)$ and (ii) $u' \in C(u)$ and $v' \in C(v)$.*

Proof. Let i be the largest integer such that $r_i < d(u,v)$. Hence $d(u,v) \leq r_{i+1} = c_1 \cdot r_i \leq (c_2 - 2) \cdot r_i$ (since c_2 was chosen to be greater than $c_1 + 2$).

Let u' and v' be the nearest nodes in the MIS V_i of $G[E_i]$ from u and v respectively. Note that u' could be u and v' could be v. Thus $d(u,u') \leq r_i$ and $d(v,v') \leq r_i$. By the triangle inequality we have, $d(u',v') \leq d(u',u) + d(u,v) + d(v,v') \leq r_i + (c_2 - 2) \cdot r_i + r_i \leq c_2 \cdot r_i < c_2 \cdot d(u,v)$. Hence, $(u',v') \in \hat{E}_i$ and also note that $d(u',v') \leq \alpha \cdot d(u,v)$ where α is any constant greater than c_2. Now note that $\{u,v\}$ is the lightest edge between a vertex in $C(u)$ and a vertex in $C(v)$ by virtue of being an MST edge. Therefore, it is the case that $u' \in C(u)$ and $v' \in C(v)$ since $d(u,u') < d(u,v)$ and $d(v,v') < d(u,v)$.

This lemma implies that for every cut (X,Y) of G and an MST edge $\{u,v\}$ that crosses the cut, there is an edge $\{u',v'\}$ in $\hat{G}$ also crossing cut (X,Y) with weight within a constant factor of the weight of $\{u,v\}$. The following result follows from this observation and properties of $\hat{E}_\ell$ proved earlier.

Theorem 2. *Algorithm 10 computes a spanning tree $\hat{\mathcal{T}}$ of G such that $wt(\hat{\mathcal{T}}) = O\left(wt\left(MST\left(G\right)\right)\right)$.*

4.3 Constant Running Time

The result of the previous subsection does not require that the underlying metric space (V, d) have constant doubling dimension. Now we assume that (V, d) has constant doubling dimension and in this setting we show that Algorithm MST-Approximation can be implemented in *constant* rounds. Even though the algorithm is described in a "sequential" style in Algorithm 10, it is easy to verify that most of the steps can be easily implemented in constant rounds in the congested clique model. However, to finish the analysis we need to show: (i) that ComputeMIS executes in constant rounds, (ii) that the $h = O(\log n)$ calls to ComputeMIS in Line 10 can be executed in parallel in constant rounds, and (iii) that MST-Sparse in Line 13 can be implemented in constant rounds. In the following, we show (iii) by simply showing that $\hat{G}$ has linear number of edges. In the previous section, we have shown (i) and later in this section we show (ii).

We first show $|\hat{E}_\ell| = O(n)$ in Lemma 19 and then argue about heavy edges.

Lemma 19. $|\hat{E}_\ell| = O(n)$.

Proof. For any edge $\{u, v\} \in \hat{E}_\ell$ either u or v or both belong to S (by construction). We orient edges such that an edge is directed towards the node in S. If both end points are in S then we add two oppositely directed edges. We prove that the out-degree of a node is bounded by a constant.
Consider a node u. Let $N_o(u)$ be the set of endpoints of all outgoing edges of u. If $|N_o(u)| < 2$ then we are done, therefore consider the case $|N_o(u)| \geq 2$. Consider any two nodes $v_i, v_j \in N_o(u)$. By construction we have, $d(u, v_i) \leq 2 \cdot d_m/n^2$ and $d(u, v_j) \leq 2 \cdot d_m/n^2$. Therefore by the triangle inequality, $d(v_i, v_j) \leq 4 \cdot d_m/n^2$. Also, by the definition of orientation $v_i, v_j \in S$ and therefore by the definition of S we have, $d(v_i, v_j) > d_m/n^2$. Hence the aspect ratio of $N_o(u)$ is at most 4. By the growth-bounded property, we have $|N_o(u)| = O(1)$. Hence, $|\hat{E}_\ell| = O(n)$.

Now we show $|\hat{E}_h| = O(n)$. We first show in the following lemma two useful properties of vertex-neighborhoods in the graph induced by $\hat{E}_i$.

Lemma 20. *For each $u \in V_i$, (i) $|N_i(u)| \leq c_3$ where $c_3 = {c_2}^{O(\rho)}$ and (ii) $N_i(u) \cup \{u\}$ induces a clique in $G[E_j]$ for all $i > 0$ and $j \geq i + \delta$ where $\delta = \left\lceil \frac{\log 2c_2}{\log c_1} \right\rceil$.*

Proof. We first show that the aspect ratio of $N_i(u)$ is bounded by $2c_2$. This follows from two facts: (a) any two points in $N_i(u)$ are at least distance r_i apart, and (b) any point in $N_i(u)$ is at distance at most $c_2 \cdot r_i$ from u and therefore, by using the triangle inequality, any two points in $N_i(u)$ are at most $2c_2 \cdot r_i$ apart. Then using the bound from the growth-bounded property we obtain the result claimed in part (i).

Now we show part (ii) of the claim. If $|N_i(u)| = 0$ then we are done. If $|N_i(u)| = 1$ then let $v \in N_i(u)$. This implies $d(u, v) \leq c_2 \cdot r_i \ < c_1^\delta \cdot r_i = r_{i+\delta}$ which implies $\{u, v\} \in E_j, j \geq i + \delta$.
Now assume $|N_i(u)| > 1$. Consider any two distinct vertices $v, w \in N_i(u)$. Since $\{u, v\}, \{u, w\} \in \hat{E}_i$ we have $d(u, v) \leq c_2 \cdot r_i$ and $d(u, w) \leq c_2 \cdot r_i$. By the triangle inequality, $d(v, w) \leq 2c_2 \cdot r_i \leq c_1^\delta \cdot r_i = c_{i+\delta}$. Therefore $\{v, w\} \in E_{i+\delta}$ and hence we have $\{v, w\} \in E_j$, for all $j \geq i + \delta$.

The implication of the above result is that $|\hat{E}_i|$ is linear in size. Since we use $O(\log n)$ layers in the algorithm, it immediately follows that $|\hat{E}_h|$ is $O(n \log n)$. However, part (ii) of the above result implies that only one of the nodes in $N_i(u)$ will be present in V_j, $j \geq i + \delta$ since V_j is an independent set of $G[E_j]$. This helps us show the sharper bound of $|\hat{E}_h| = O(n)$ in the following.

Without loss of generality assume that h is a multiple of δ (if not, add at most $\delta - 1$ empty layers $\hat{E}_{h+1}, \hat{E}_{h+2}, \ldots$ to ensure that this is the case). Let

$$\beta(j) = \bigcup_{i=(j-1)\delta+1}^{j\delta} \hat{E}_i \qquad \text{for } j = 1, 2, \ldots, \frac{h}{\delta}$$

be a partition of the layers $\hat{E}_i$ into *bands* of δ consecutive layers. Let $\hat{E}_{odd} = \cup_{j:odd} \beta(j)$ and $\hat{E}_{even} = \cup_{j:even} \beta(j)$.

Lemma 21. $|\hat{E}_{odd}| = O(n)$, $|\hat{E}_{even}| = O(n)$ *and therefore* $|\hat{E}| = O(n)$.

Proof. We prove the claim for $\hat{E}_{odd}$. The proof is essentially the same for $\hat{E}_{even}$. We aim to prove the following claim by induction on k (for odd k): for some constant $C > 0$,

$$\left| \bigcup_{j:odd \geq k} \beta(j) \right| \leq C \cdot \left| \bigcup_{j:odd \geq k} V(j) \right|, \tag{2}$$

where $V(j)$ is the set of vertices such that every vertex in $V(j)$ has some incident edge in $\beta(j)$. Setting $k = 1$ in the above inequality, we see that $|\hat{E}_{odd}| = |\cup_{j:odd \geq k} \beta(j)| = O(n)$. To prove the base case, let k' be the largest odd integer less than or equal to h/δ. Then, $\cup_{j:odd \geq k'} \beta(j) = \beta(k')$ and $\cup_{j:odd \geq k'} V(j) = V(k')$. Consider a vertex $v \in V(k')$. By Lemma 20, there are at most c_3 edges incident on v from any layer. There are δ layers in $\beta(k')$ and therefore there are at most $c_3\delta$ edges from $\beta(k')$ incident on any vertex $v \in V(k')$. Hence, $|\beta(k')| \leq c_3\delta|V(k')|$. Therefore, for any constant $C \geq c_3\delta$, it is the case that $|\cup_{j \geq k'} \beta(j)| \leq C \cdot |\cup_{j \geq k'} V(j)|$.

Taking (2) to be the inductive hypothesis, let us now consider $|\cup_{j \geq k-2} \beta(j)|$. Then,

$$\left| \bigcup_{j:odd \geq k-2} \beta(j) \right| \leq \left| \bigcup_{j:odd \geq k} \beta(j) \right| + |\beta(k-2)| \leq C \cdot \left| \bigcup_{j:odd \geq k} V(j) \right| + c_3\delta \cdot |V(k-2)|. \tag{3}$$

The second inequality is obtained by applying the inductive hypothesis and the inequality $|\beta(k-2)| \leq c_3\delta|V(k-2)|$. By Lemma 20, at most half the vertices in $V(k-2)$ appear in $\cup_{j \geq k} V(k)$. Therefore, $|V(k-2) \setminus (\cup_{j \geq k} V(j))| \geq |V(k-2)|/2$. Hence,

$$\left| \bigcup_{j:odd \geq k-2} \beta(j) \right| \leq C \cdot \left| \bigcup_{j:odd \geq k} V(j) \right| + 2c_3\delta \cdot \left| V(k-2) \setminus \left(\bigcup_{j:odd \geq k} V(j) \right) \right|.$$

Picking $C \geq 2c_3\delta$, we then see that

$$\left| \bigcup_{j:odd \geq k-2} \beta(j) \right| \leq C \cdot \left(\left| \bigcup_{j:odd \geq k} V(j) \right| + \left| V(k-2) \setminus \left(\bigcup_{j:odd \geq k} V(j) \right) \right| \right) = C \cdot \left| \bigcup_{j:odd \geq k-2} V(j) \right|.$$

The result follows by induction.

4.4 Many MIS Computations in Parallel

In this section, we argue that Algorithm 6 LowDimensionMIS can be executed on the $O(\log n)$ different distance threshold graphs in parallel on a congested clique. Table 1 shows number of messages sent/received per node in the execution of Algorithm 6 and from this it is easy to see that Line 8 of Phase 2 can be executed as it is using Lenzen's routing protocol in $O(1)$ rounds for all the $O(\log n)$ layers in parallel due to their low communication requirements. For Lines 4-6 of Phase 2 we do the following load balancing via a *designated receiver scheme*: each vertex has to send at most $O(n^{1/4} \log n)$ messages in an execution of Phase 2 for a layer. Therefore, for $O(\log n)$ layers one node is responsible of sending $O(n^{1/4} \log^2 n)$ messages. There are only $\lceil 2 \log n \rceil$ receivers needed for in an execution at a layer. For all layers the number of receivers needed are $O(\log^2 n)$. Hence we can designate different receivers such that no receiver gets more than $O(n)$ messages in execution of Phase 2 for all layers. Similar designated receiver scheme is applied for the execution of Phase 1.

For parallel execution of Line 9 (SequentialMIS) of Phase 4 for all $O(\log n)$ layers we use the following *message encoding scheme*: Each vertex v constructs a $O(\log n)$-length bit string specifying 1 at position ℓ if v is in MIS for the layer ℓ otherwise 0. Each vertex v broadcasts this string. For a layer ℓ, each vertex considers only ℓ^{th} bit of this message.

5 Constant-Approximation to MFL

Berns et al. [3,2] showed how to compute a constant-factor approximation to MFL in expected $O(\log \log n)$ rounds. (The algorithm presented in [2] runs in expected $O(\log \log n \cdot \log^* n)$ rounds, but this was subsequently

Table 1. Number of messages sent/received per node in the execution of Algorithm 6

Phase	Line	Analysis	Number of messages to send per node	Number of receivers	Number of messages to receive per receiver
1	2-4	Lemma 8	$O(n^{1/2})$	$n^{1/2}$	$O(n)$
2	4-6	Lemma 11	$O(n^{1/4}\log n)$	$\lceil 2\log n \rceil$	$O(n)$
	8	Lemma 12	$O\left(\text{poly}(\log n)\right)$	n	$O\left(\text{poly}(\log n)\right)$
3	-	Thm. 1	$O(n^{1/2}\,\text{poly}(\log^* n))$	n	$O(n^{1/2}\,\text{poly}(\log^* n))$
4	3	Lemma 14	$O(1)$	1	$O(n)$
	9	Lemma 14	1 (1-bit)	n	n

improved to expected $O(\log\log n)$ in [3].) A high level description of this algorithm is as follows. Each node v locally computes a value $r_v \geq 0$ that is a function of its opening cost f_v and distances to other nodes $\{d(v,w) \mid w \in V\}$. Nodes with similar r_v-values join the same class; more precisely, a node v with $3^k \cdot r_m \leq r_v \leq 3^{k+1} \cdot r_m$, joins a class V_k. Here r_m is the minimum r_u-value over all nodes $u \in V$. For nodes in each class V_k, we construct a graph $H_k = (V_k, E_k)$, where the edge-set E_k is defined as $\{\{u,v\} \mid u,v \in V_k, d(u,v) \leq r_u + r_v\}$. In the rest of the algorithm, in order to figure out which nodes to open as facilities, the algorithm computes a t-ruling set on each graph G_k. Analysis in [3,2] then shows that the solution to facility location produced by this algorithm is an $O(t)$-approximation. In [3] it is shown how to compute a 2-ruling set in expected $O(\log\log n)$ rounds on a congested clique. Since the classes V_k form a partition of the nodes, the ruling set computations occur on disjoint sets of nodes and can proceed in parallel. This leads to a constant-factor approximation to MFL in expected $O(\log\log n)$ rounds.

The 3-ruling set algorithm and the MIS algorithm in the present paper can replace the slower 2-ruling set and this yields the following result.

Theorem 3. *There exists a distributed algorithm that computes a constant-approximation to the metric facility location problem (w.h.p.) in the congested-clique model and which has an expected running time of $O(\log\log\log n)$ rounds. Additionally, if the input metric space has constant doubling dimension then a constant-approximation can be computed in constant rounds (w.h.p.)*

6 Conclusion

In a recent paper, Drucker et al. [7] show that the congested clique can simulate powerful classes of bounded-depth circuits, implying that even slightly super-constant lower bounds for the congested clique would give new lower bounds in circuit complexity. This provides some explanation for why there are no non-trivial lower bounds in the congested clique model. One could view this result as providing motivation for proving even stronger upper bounds. As shown in this paper, it is possible to design algorithms that run significantly faster that $\Theta(\log\log n)$ rounds for well-known problems. Continuing this program, we are interested in designing algorithms running in $o(\log\log n)$ rounds for MST and related problems such as connectivity verification.

Acknowledgments. We would like to thank reviewers of DISC 2014 for their careful reading and thoughtful comments.

References

1. M. L. Balinski. On finding integer solutions to linear programs. In *Proceedings of IBM Scientific Computing Symposium on Combinatorial Problems*, pages 225–248, 1966.
2. Andrew Berns, James Hegeman, and Sriram V. Pemmaraju. Super-Fast Distributed Algorithms for Metric Facility Location. In Artur Czumaj, Kurt Mehlhorn, Andrew M. Pitts, and Roger Wattenhofer, editors, *ICALP (2)* , volume 7392 of *Lecture Notes in Computer Science*, pages 428–439. Springer, 2012.
3. Andrew Berns, James Hegeman, and Sriram V. Pemmaraju. Super-Fast Distributed Algorithms for Metric Facility Location. *CoRR*, abs/1308.2473, August 2013.
4. G. Cornuejols, G. Nemhouser, and L. Wolsey. *Discrete Location Theory.* Wiley, 1990.

5. Mirela Damian, Saurav Pandit, and Sriram V. Pemmaraju. Distributed Spanner Construction in Doubling Metric Spaces. In *International Conference on Principles of Distributed Systems*, volume 4305 of *OPODIS*, pages 157–171. Springer, 2006.
6. Jeffrey Dean and Sanjay Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. *Commun. ACM*, 51(1):107–113, January 2008.
7. Andrew Drucker, Fabian Kuhn, and Rotem Oshman. On the power of the congested clique model. In *Proceedings of the 2014 ACM Symposium on Principles of Distributed Computing*, PODC '14, pages 367–376, New York, NY, USA, 2014. ACM.
8. Devdatt P. Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009.
9. M. V. Eede, P. Hansen, and L. Kaufman. A plant and warehouse location problem. *Operational Research Quarterly*, 28(3):547——554, 1977.
10. C. Frank. *Algorithms for Sensor and Ad Hoc Networks*. Springer, 2007.
11. Joachim Gehweiler, Christiane Lammersen, and Christian Sohler. A Distributed O(1)-approximation Algorithm for the Uniform Facility Location Problem. In *Proceedings of the Eighteenth Annual ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '06, pages 237–243, New York, NY, USA, 2006. ACM.
12. James Hegeman and Sriram V. Pemmaraju. A Super-Fast Distributed Algorithm for Bipartite Metric Facility Location. In *DISC*, pages 522–536, 2013.
13. James W. Hegeman and Sriram V. Pemmaraju. Lessons from the congested clique applied to mapreduce. In *SIROCCO*, SIROCCO'14, pages 149–164, 2014.
14. Howard Karloff, Siddharth Suri, and Sergei Vassilvitskii. A model of computation for mapreduce. In *Proceedings of the Twenty-first Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '10, pages 938–948, Philadelphia, PA, USA, 2010. Society for Industrial and Applied Mathematics.
15. Hartmut Klauck, Danupon Nanongkai, Gopal Pandurangan, and Peter Robinson. The Distributed Complexity of Large-scale Graph Processing. *CoRR*, abs/1311.6209, 2013.
16. Fabian Kuhn, Thomas Moscibroda, and Rogert Wattenhofer. On the Locality of Bounded Growth. In *Proceedings of the Twenty-fourth Annual ACM Symposium on Principles of Distributed Computing*, PODC '05, pages 60–68, New York, NY, USA, 2005. ACM.
17. Christoph Lenzen. Optimal Deterministic Routing and Sorting on the Congested Clique. In *Proceedings of the 2013 ACM Symposium on Principles of Distributed Computing*, PODC '13, pages 42–50, 2013.
18. Christoph Lenzen and Roger Wattenhofer. Brief announcement: exponential speed-up of local algorithms using non-local communication. In *Proceedings of the 2010 ACM Symposium on Principles of Distributed Computing*, PODC '10, pages 295–296, 2010.
19. Zvi Lotker, Boaz Patt-Shamir, and David Peleg. Distributed MST for Constant Diameter Graphs. *Distributed Computing*, 18(6):453–460, 2006.
20. Grzegorz Malewicz, Matthew H. Austern, Aart J.C Bik, James C. Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. Pregel: A System for Large-scale Graph Processing. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*, SIGMOD '10, pages 135–146, New York, NY, USA, 2010. ACM.
21. Thomas Moscibroda and Roger Wattenhofer. Facility location: distributed approximation. In *PODC '05: Proceedings of the twenty-fourth annual ACM symposium on Principles of distributed computing*, pages 108–117, 2005.
22. S. Pandit and S. V. Pemmaraju. Finding facilities fast. *Distributed Computing and Networking*, pages 11——24, 2009.
23. David Peleg. *Distributed Computing: A Locality-Sensitive Approach*, volume 5. Society for Industrial Mathematics, 2000.
24. Johannes Schneider and Roger Wattenhofer. A Log-Star Distributed Maximal Independent Set Algorithm for Growth-Bounded Graphs. In *Proceedings of the twenty-seventh ACM symposium on Principles of distributed computing*, pages 35–44. ACM, 2008.